The Effect of Neutral Density Estimation Errors on Satellite Conjunction Serious Event Rates

M.D. Hejduk[1] and D.E. Snow[2]

[1]Astrorum Consulting LLC, in Support of NASA Goddard Space Flight Center.

[2]HQ Air Force Space Command Studies and Analysis Division (A2/3/6Z).

Corresponding author: Matthew Hejduk (mdhejduk@astrorum.us)

Key Points:

- Atmospheric modeling effects on satellite conjunction assessment are best assessed by the effect on high-risk conjunction event rates
- Including density model errors in conjunction risk calculations greatly improves the stability and durability of the results
- Density model accuracy improvements will improve conjunction assessment but only if model errors are characterized.

Abstract

While past studies have investigated the effect of neutral atmospheric density mismodeling on satellite conjunction (CA) assessment, none has focused their investigation specifically on serious (high-risk) conjunction events, which are the event types that drive both risk and workload for CA operations. The present study seeks to do this by reprocessing a large number of archived actual conjunction events, artificially introducing atmospheric density error to these events, and then examining the effect of these introduced errors on the probability of collision (Pc) calculation, which is the principal parameter used to assess collision risk. These reprocessed calculations are executed both with the satellites' covariances unaltered and with a covariance modification that accounts for the induced atmospheric density error. The results indicate that the situation is greatly aided by an *a priori* knowledge of the approximate density estimation error, even if the model itself is unaltered—missed detections (Type II errors) due to density estimation uncertainty are substantially reduced when the density model prediction error is characterized and can be included in the satellite covariance and thus Pc calculation. Overall improvements in density model predictive performance are also important to improving CA, especially for false alarm (Type I error) reduction; but model enhancements that include a robust, in-model error analysis offer the most significant improvements to the CA enterprise.

1 Introduction

Conjunction Assessment (CA), or the determination of the possibility and likelihood of collision between space objects, has become an area of greatly increased interest and research within space situational awareness. The increase in satellite population, the deployment of sensors that can track smaller and smaller objects, and United States Strategic Command's more expansive data-release policies to commercial and foreign entities have all combined to create a significantly expanded CA solution space: both an increased number of known conjunctions between space objects, and the distribution of relevant conjunction data to the concerned entities so that further analysis and potential conjunction remediation can be performed. Considerable academic research has accompanied this increased operational interest and been directed to all aspects of the problem, from the basic astrodynamics of identifying conjunctions, to the determination of collision likelihood, to efficient methods for calculating conjunction remediation options. In analyzing the entire chain of CA-related calculations for potential sources of error, drag acceleration error suggests itself immediately as the single greatest contributor to satellite state propagation uncertainty (and thus CA calculation error) for satellites with altitudes less than about 1000 km; one of the principal ingredients to the drag acceleration calculation is neutral atmospheric density estimation. While substantial literature exists on comparative performance among different atmospheric density models when presented with different space weather phenomena, two studies focused specifically on the relationship between density mismodeling and resultant satellite position errors (Emmert *et al.* 2016; Vallado and Finkleman 2014) have shown that mismodeling effects can be substantial, potentially large enough to influence CA-related operational conclusions. Emmert (2014) works out some bounding calculations to try to quantify this effect, using an idealized conjunction assessment framework.

The purpose of the present work is to enhance the efforts referenced above by quantifying this effect within the CA calculation, risk assessment, and decision support methodologies that follow the best practices of the leading CA institutions. The NASA Robotic Conjunction

Assessment Risk Analysis (CARA) project is the largest single governmental organization to conduct routine CA operations. With approximately 70 protected primary spacecraft spanning a wide variety of orbit regimes and a full-time research staff, CARA has been active nationally and internationally in developing, operationally deploying, and evaluating CA best practices. It will be against these practices that atmospheric density error will be examined to determine its effect on the number of serious conjunctions and therefore actual imputed CA workload. The result of such an investigation can be used to construct a more definitive statement of the operational impact of atmospheric density error and therefore make evident the operational benefits of improved atmospheric modeling.

This study thus follows the following organizational schema. First, basics of the CA enterprise will be discussed in order to establish the proper framework for subsequent discussions. Next, the relationship between errors in atmospheric drag acceleration and the calculation of the probability of collision (Pc), the basic parameter for performing conjunction risk assessment, are delineated in order to establish the link between atmospheric density estimation errors and changes in the assessed risk of a particular conjunction and thus formulate the study's first key question: how strongly are Pc values affected by atmospheric density modeling errors. After this, methods for attempting to compensate for these density estimation errors within a satellite's state covariance matrix are discussed, and with this the investigation's second key question is defined: by how much are CA risk assessment errors that arise from atmospheric density estimation errors attenuated by properly adjusting the satellite's covariance matrix to account for these errors. With these background items presented, one is then ready to describe the dataset and conduct of the experiment itself and its results. Finally, research objectives that stem from these results are sketched out and helpful additional future work is summarized.

2 The CA Enterprise and Associated Processes

The CA enterprise is typically divided into three parts. While some of these divisions are not intrinsic to the calculations themselves, they are nonetheless both conceptual and practical divisions guiding the way that conjunction-related data are presently generated and distributed to the agencies responsible for risk assessment; they thus serve as useful points of demarcation of the process.

2.1 CA Screenings

The purpose of CA screenings is to discover potential conjunctions between space objects some time in advance of the time of closest approach (TCA) so that risk assessment activities can be performed, orbit determination (OD) refinements can be executed, and, if necessary, conjunction remediation actions taken. A screening is an evaluation of a single protected asset's future positions in comparison to the future positions of all other objects in the space catalogue. Typically performed for a look-ahead period of at least seven days, the ephemeris of the protected object (called the "primary") is compared to ephemerides for all the other catalogued objects (called "secondaries"); and any secondaries that come within a specified (componentized) distance of the primary are identified as possible conjunctions, which can then be sent on to the risk assessment portion of the process. To be specific, a particular volumetric region is constructed about the primary object and "flown" along its trajectory; any penetration of that volume by a secondary object constitutes a conjunction. These screening volumes are

rarely spherical but usually ellipsoidal in order to align with the expected satellite state error distribution in the particular orbit regime (Narvet *et al.*, 2011). A number of efficient filtering mechanisms have been developed in order to decrease the computational burden of this process (*e.g.,* Hoots *et al.* 1984, George and Chan 2012, Alfano 2013).

2.2 CA Risk Assessment

The identification of a conjunction between a protected primary and a secondary object is not a datum without utility, but from a satellite protection point of view it alone is not actionable information. In an earlier period of CA operations, risk assessments were attempted based solely on the predicted closest miss distance between the primary and secondary. However, as the discipline matured, this methodology was found to be less than fully adequate because it did not consider the state uncertainties between the two objects. If the state uncertainties are large, a small miss distance does not indicate a high likelihood of a collision because the actual satellites' positions could well be far from the mean values. Similarly, a miss distance that does not seem particularly small, if it aligns properly with the actual state uncertainty values, can produce a situation in which the collision likelihood is larger than one might expect. It thus became clear that a collision likelihood, or probability of collision (Pc), calculation should be developed in order to provide a statement of actual collision risk. The initial assembly of a Pc calculation methodology was performed for the Space Shuttle program (Foster and Estes 1992), and since that time a number of calculation approaches have been developed and capably summarized by Chan (2008). The availability of precision catalogue and accompanying covariance data has allowed these probabilistic calculations to be performed by nearly all CA practitioners. A full risk assessment evaluation includes not only the Pc but also an assessment of the adequacy of the primary and secondary OD, a consideration of the uncertainty of the Pc itself, and prognostications of what is likely to take place regarding additional secondary tracking and therefore at what particular time a final risk assessment calculation might be best made (Newman *et al.* 2014, Hejduk and Johnson 2016).

2.3 CA Serious Event Remediation

If the collision risk is considered high, usually because the Pc value exceeds a particular threshold and the OD results appear credible, the focus of the process turns to conjunction remediation. This is typically accomplished through the execution of a specially-assembled satellite maneuver, called a Risk Mitigation Maneuver (RMM), to change the primary satellite's trajectory in order to reduce the risk of the conjunction to an acceptable level. Sometimes a similar outcome can be achieved by changing the time and/or intensity of a maneuver already scheduled for a different purpose (such as a drag make-up maneuver) in a manner that both accomplishes the maneuver's original objectives and also mitigates the conjunction risk. The planning process typically is a nested analysis of the resultant Pc for the main conjunction versus potential instantaneous burn intensity and burn time. The basic trade-off is that burns conducted earlier (*i.e.,* longer before TCA) can be smaller yet produce the same remediative value; but waiting longer increases the likelihood that additional satellite tracking will refine the orbits such that the collision risk drops considerably, thus obviating the need for a maneuver at all. Satellites that use longer burns, electric propulsion, or other methods for trajectory control such as differential drag, require modified remediation methods. However, the overall concept and procedural outline is identical.

The focus of the present analysis is on the risk assessment stage of the process: to determine the effect of atmospheric density mismodeling on the number of CA serious events; for it is these events that drive nearly all of the CA analytical and communications workload. The previously-referenced works by Emmert *et al.*, in which he develops analytical expressions that link density mismodeling error and satellite position error and in which he calculates expected differences in the number of conjunctions found by screening processes, are both excellent studies; but they do not formulate conclusions within the framework of current CA risk assessment practices and therefore cannot be immediately deployed to assess the expected effect of mismodeling on CA operations. To examine this issue so that operational conclusions can be drawn, two things are necessary: first, actual conjunctions with typically-encountered conjunction geometries and Pc calculations must be analyzed; and second, the effects of density mismodeling on state estimate covariances, including common approaches to try to account for such modeling errors, must be included. The next two sections provide a more extended treatment of the relationship of density mismodeling to Pc calculations and discuss atmospheric error compensation strategies often applied to state covariances.

3 Relationship of Drag Acceleration Error to Resultant Pc

Using Montebruck and Gill's (2005) notation, the satellite acceleration due to drag is given as

$$\ddot{r} = -\frac{1}{2} C_D \frac{A}{M} \rho v_r^2 e_v \ , \tag{1}$$

in which *r*-double-dot is the anti-velocity acceleration, C_D is the drag coefficient (dimensionless), A is the spacecraft frontal area (normal to the velocity vector), M is the spacecraft mass, ρ is the atmospheric density, v_r is the magnitude of the velocity relative to the atmosphere, and e_v is the unit vector in the direction of the spacecraft velocity. Because the spacecraft mass and frontal area are generally not known, the group quantity of C_DA/M is solved for as a unit and called the ballistic coefficient. Two conclusions can be drawn immediately from the formulation in (1). First, the atmospheric density is multiplicatively linked to all other terms in the drag calculation, so any errors in the density estimate flow directly to the calculated acceleration: if the density estimate is incorrect by a factor of 50%, the resultant drag acceleration calculation will be misrepresented by that same amount. Second, because all of the terms in the expression are multiplicatively combined, the ballistic coefficient and the density estimate can be used as aliases for each other; if one wishes to increase or decrease the atmospheric density by a certain percentage, this can be accomplished by varying the ballistic coefficient by the same percentage.

Because it acts in the anti-velocity direction, the immediate and largest manifestation of drag acceleration, and similarly drag acceleration error, is in the satellite in-track velocity and therefore in-track position. This change in in-track satellite velocity also has a secondary effect on the orbit's semi-major axis and thus the satellite radial position and velocity. For most conjunctions, radial separation between the two orbits tends to govern the Pc calculation the most strongly; so errors that affect the radial component are quite likely to affect the satellites' relative position at TCA and therefore the collision probability.

However, it must be remembered that the covariance contributes substantially to the Pc calculation, and adjustments to the covariance to account for atmospheric density mismodeling can allow a correct and useful Pc to be calculated even if the induced position estimation errors

themselves cannot be remediated. The Pc represents the likelihood that, given the uncertainties in the two satellites' positions at TCA, their actual miss distance will be smaller than a specified tolerance, called the hard-body radius (HBR); this would be considered the equivalent of a collision (two satellite's flying closer than the HBR means that the "hard bodies" of the two satellites could contact each other, although it is possible that due to fortunate alignment of the two vehicles an actual collision may not take place). If the uncertainty in the density estimates were characterized and known, there are techniques to include this uncertainty in the covariances and therefore the Pc calculation. This allows a probability calculation to be executed that accurately represents the risks, given all of the known uncertainties of the situation.

The effect of increasing the state estimate uncertainty does not have a monotonic effect on the resultant Pc; rather, it is a function of the ratio of the size of the joint uncertainty of the two objects' state estimates to the miss distance. This is a subtle but important point and thus bears some additional discussion, aided by Figure 1 below, which provides a curve that gives the Pc value (as a ratio to its maximum value) as a function of the ratio of joint covariance size to miss distance (for simplicity, a spherical covariance is used here, with its size represented by the sphere's radius). So long as the analysis is limited to situations with reasonable miss distances (*i.e.*, not ridiculously large or vanishingly small), one can see that very large and very small combined covariances push the Pc to very small values, with a peak in the middle. In thinking about the implications of these conditions, this result makes physical sense. It is important to recall that OD processes produce an estimate of a mean state and a covariance that indicates the expected position (and velocity) dispersion about that state. When the uncertainties are very large, an estimate of the mean state is still the expected value, but it is not a very strong expression of central tendency; in the limit as the variance moves to infinity, the mean becomes merely the center point in a uniform distribution of vanishingly small density (at least for error in a single position component, which can be expected to follow a Gaussian distribution). So while the mean is still the most likely value, the uncertainty is so large that the likelihood of the two objects being at the mean position indicated by the OD and propagation processes is not great and the Pc is small. Thus, if the two satellites' positions are poorly determined, then on the basis of the information available, the strength of a conclusion that they will actually pass in close proximity of each other will have to be weak. Conversely, when the uncertainties are very small, if the miss distance at TCA is notably greater than the HBR, one can conclude that the likelihood of an actual collision is low, thus also producing a small Pc. One can thus achieve a small Pc for two different reasons: because the knowledge of the satellites' states is poor enough that a definitive conclusion of a dangerously close approach cannot be credibly extracted from the data, or because the knowledge of the satellites' states is so good that one can state with certainty that a dangerously close approach will not occur. Alfano (2005) pointed out the difference between these two ways of achieving a low Pc and argued preferentially for the latter, as it is a conclusion stemming from good data rather than poorer data; Frisbee (2009) noted in response that in either case the Pc is an appropriate assessment of the collision risk and that, while one would always prefer better to poorer data to enable decision-making, there is nothing improper or illegitimate in using the Pc from either "side" of the curve in Figure 1 for operational CA decisions.

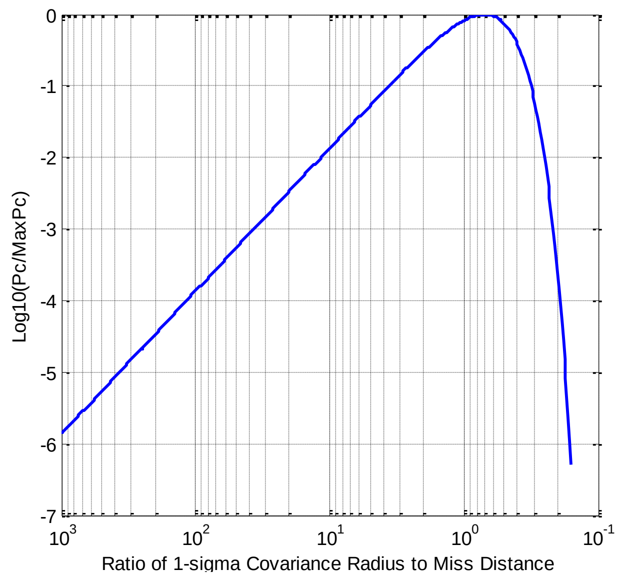

Figure. 1: Relative Pc Plot as a function of ratio of covariance size to miss distance

Adjusting the covariance to account for atmospheric density estimation error aligns with the concept of reflecting all known state estimation errors into the accompanying state covariance matrix. While including this additional error source in the covariance can be expected to make the covariance larger, this increase, as discussed above, will not in all cases decrease the Pc value. The effect will depend on the ratio of the covariance size to the miss distance before the atmospheric density estimation error compensation is added. If one is well on the right side of the above curve, a modest increase in covariance size can have the effect of increasing the Pc substantially and thus changing the characterization of the event from "dismissible" to "serious." Extremely large density errors will probably decrease the resultant Pc in most cases, but that outcome is not necessarily undesirable: if the uncertainty is properly represented, then the Pc reflects the actual likelihood of collision given the quality of the data available at the time an operational decision is to be rendered. The next section outlines some of the proposed and deployed methods for adjusting the covariance to accommodate this known source of error.

4 Accounting for Neutral Density Mismodeling within State Covariances

Reflecting neutral density estimation error within state covariance matrices is a subset of the general problem of covariance or uncertainty realism, a topic that has been treated at length in a recent report of the Air Force Space Command Astrodynamics Innovation Committee (Poore *et al.*, 2016). There are a number of techniques that are frequently employed to improve the realism of covariances—that is, their ability realistically and reliably to represent the actual state errors. Each of the major techniques is discussed briefly below.

4.1 Covariance Scaling

The governing presumption in single-factor covariance scaling is that the covariance is over- or under-sized by a scalar amount, so the covariance can simply be multiplied by a factor

(actually typically the square of a factor) to make it more representative. Such factors are determined by analyzing past precision orbit data on the satellite in question (or ensemble sets of satellites with similar orbit maintenance properties), examining the relationship between the actual state errors and the statistical error summary represented by the covariance, and calculating a scale factor that will size the covariance so that it may accurately represent the statistical distribution of the actual empirical state errors (Hejduk 2013, Poore *et al.* 2016). A natural extension of this approach is to calculate a scale factor for each of the three position components; if S represents a vector of three scale factors, scaling is accomplished by the matrix multiplication $S*C*S^T$. This approach is generally used to perform omnibus covariance realism corrections rather than to respond to a particular source of error (such as that from neutral density mismodeling).

4.2 Covariance Correction Matrices

An additional level of sophistication to simple covariance scaling was introduced through the work of Cerven (2011, 2013). Here, entire correction matrices, rather than simple scaling vectors, are produced in order to attempt to correct the matrix for both size and orientation. Pre- and post-multiplying the covariance by the correction matrix produces a properly sized and oriented covariance. Again, applications of Cerven's method to this point have attempted only an ensemble correction for all covariance realism errors rather than those arising from a single source.

4.3 Physically-Connected Process Noise

When using one of the many strains of Kalman filters to propagate covariances, provision is usually made to include process noise, or characterized acceleration uncertainties, which grow over time and can be propagated along with the covariance in order to increase the covariance size appropriately. In astrodynamics this was originally deployed to reflect characterized uncertainties in the applied geopotential model, for which processing requirements imposed practical limitations on the order of the model that could be applied; but in more recent times it has been used as an omnibus covariance correction methodology (Vallado *et al.*, 2010) and as part of dedicated covariance realism efforts (Duncan and Long 2006; Zaidi and Hejduk 2016).

4.4 Consider Covariance Parameters

A traditional method to adjust covariances to make them more representative is through the use of consider parameters (Tapley, Schutz, and Born 2004). Such adjustments are determined from *a priori* error information external to the OD and are thus not solved-for but rather "considered" as part of the estimation process. This methodology has been embraced by US Strategic Command, the Department of Defense (DoD) entity that supervises the production and distribution of the space catalogue; and its purpose is to compensate for expected state propagation errors due principally to atmospheric density mismodeling. The covariance as formulated does not contain an atmospheric density uncertainty term directly; but because atmospheric density and the ballistic coefficient are multiplicatively coupled, one can alter the ballistic coefficient variance, which does appear in the covariance, and achieve an equivalent effect. If the percent error of the atmospheric density estimation is known, then the square of this amount can be added directly to the ballistic coefficient variance. As the covariance is propagated (the usual method is through pre-and post-multiplication by a state transition matrix), this increased variance will manifest an effect on the propagated covariance's position and

velocity variances, as well as cross-correlation terms. In the present study, it is this approach (because it aligns with current DoD operational practice) that will be used to attempt to account for characterized atmospheric density modeling error.

All of these methods do improve the situation because they allow the covariance to accommodate density errors when they are known. Of course, the actual state errors due to density mismodeling are still present; and for the CA problem, in most cases this will cause the miss distance between the two satellites to be incorrectly calculated, affecting the rectitude of the Pc. The only complete solution is to develop and deploy an atmospheric model with very low prediction errors. Nonetheless, the ability to incorporate error into the covariance is a preferable position over not considering the error at all, as it properly folds this uncertainty into the probabilistic calculation.

5 Experiment to determine Effects of Neutral Density Mismodeling on CA Calculations

5.1 Categorization Method for Severity of Conjunction Events

As stated previously, while investigations of effects on CA screening results or satellite propagated position errors are of interest, it is changes to the number of serious conjunction events that actually affect CA operational decisions and workload. Therfore, it is this parameter that must be determined in order to ascertain the actual effect of atmospheric mismodeling on the CA enterprise. It is straightforward to reprocess historical conjunctions using different atmospheric density values in order to see the effect on the resultant Pc, and in fact such an approach is the basis for the present study; but simple comparison of ranges of Pc values gives data but not information that is properly operationally contextualized and therefore not particularly easy to interpret definitively. Instead, it has been shown to be helpful operationally to view Pc data within a “color coding” framework similar to one used by the CARA project as its method to communicate conjunction severity (Newman *et al.*, 2014), and this approach is explained below.

Green conjunctions are conjunctions with a calculated Pc < 1E-07. Such conjunctions will rarely rise to a level at which they would be considered serious, so no additional analysis effort is directed to them. Should the Pc increase as the event develops, the event can be recategorized. However, this particular threshold value was chosen so that the number of green events that eventually become serious remains below 0.1%

Yellow conjunctions are conjunctions with a calculated Pc between 1E-07 and ~1E-04. Conjunctions in this range do have a reasonable likelihood of becoming serious as each event develops, so they are given additional monitoring, which typically includes manual inspection and massaging of the object ODs and requests for additional tracking, if it is felt that supplemental tracking will improve the OD and therefore the Pc calculation. The yellow category is typically a way-station for conjunctions on a path to either a red or green status (usually the latter), but some number of events reach their TCA still in the yellow category.

Red events are conjunctions whose calculated Pc is 1E-04 – 4E-04 and higher (the precise value depends to some degree on the particular protected spacecraft; for the purposes of the present study, a threshold value of 1E-04 has been used). Such events are considered serious, and their presence engenders a considerable amount of additional analysis, including the construction of a High Interest Event presentation to be delivered in person to the satellite owner/operator. Events still in red status at about three days to TCA will often require

remediation. At that point, a parallel analytical effort is also begun to examine conjunction remediation options based on the methods available to the particular spacecraft and to make a series of recommendations to the owner/operator.

Given this approach to event categorization, the present analysis can usefully frame results in terms of the color changes sustained by events as atmospheric density estimation error is added to the Pc calculation; this will reveal the sensitivity of the Pc calculation to such density errors. For example, if a particular event's color, with density unaltered, is green, the density can be increased/decreased incrementally up to 100% and the color change of the event, if any, noted for each increase/decrease level. This change can be done both with and without attempted modifications to the covariance to account for expected atmospheric density mismodeling error. These color changes can be thought of more classically as both Type I (false alarm) and Type II (missed detection) errors. Events with perfect density knowledge would be red or yellow but with density error become yellow or green are Type II errors and are certainly the more worrisome—these are either serious or concerning conjunctions that, due to density modeling errors, are not being afforded the requisite amount of analysis and attention and not being considered for potential remediation. However, given the workload associated with non-green events, Type I errors (here events that are yellow or green when perfect atmospheric density values are used but instead emerge as red or yellow when errors are introduced) are also of concern. Significant amounts of analysis and workload, as well as risk mitigation orbit alterations and all of the inherent risks that those processes contain, can be applied to events that in actuality are not risky. So it is operationally important to try to reduce both types of errors as much as possible.

5.2 Proposed Dataset and Experiment to Determine CA Effects of Neutral Density Mismodeling

The sensitivity of conjunctions to atmospheric density mismodeling can be characterized by reprocessing historical events in the following way. The nominal processing of the event, namely propagating both objects' states and covariances forward to TCA and subsequently calculating the Pc, can be considered the truth datum for the event. It does not matter whether this result actually represents the results of a high-fidelity density model or historically-accurate reconstructed density values; for what is of interest here is how modifying the density values from this baseline calculation, which gives the effect of adding density estimation error, changes the Pc calculation. Density error can then be added and the satellite states/covariances and Pc recalculated. These recalculated values can be expressed in terms of whether the event color changed from the color associated with the "truth" version of the event. An additional parallel calculation with each density alteration event is to fold the density error into the covariance as a consider parameter addition to the ballistic coefficient variance (as described in section 4.4 above). This emulates the situation in which the density estimation error has been appropriately characterized beforehand and can be accounted for in the state uncertainty representation. Of interest will be the degree to which the including of characterized error in the covariance generates an event color different from that in the situation in which the density error is not included in the covariance in this way.

The particulars of the conducted experiment are as follows. Three days in the recent past for which there was moderate solar activity (Ap values of 10 to 25; one day in April 2015, November 2015, and May 2016) were identified. This condition for solar activity was probably

not strictly necessary but was chosen in order to give reasonable exercise of the atmospheric model (as opposed to a completely quiescent period) without imposing a truly difficult-to-model situation, such as an unfolding solar storm. All of the CARA conjunctions observed for these three days were collected and winnowed down to include only those against protected assets with a perigee height of less than 1000 km. Each event was then reprocessed with a series of global alterations to the density value, ranging from -100% to +100%. Specifically, the following density scale factors were used: [1 1.1 1.2 1.25 1.3 1.4 1.5 1.75 2] and their reciprocals. As discussed in an earlier section, it is possible to effect a multiplicative change to the density value by modifying the ballistic coefficient by the same scale factor. This method for density adjustment was used in the present analysis. In each case for which which the density has been adjusted, the propagation and Pc computation were executed, both with the covariance unaltered and with the covariance's ballistic coefficient variance additively increased by the square of the expected percent density error, to emulate the situation in which a reasonable estimate of the atmospheric model error is known at the time of computation. Each conjunction was thus reprocessed 32 times (16 different scale factors [excluding unity], with both a compensated and uncompensated covariance for each).

6 Experiment Results

Figures 2-4 present the experiment's results, and each figure is a collection of three plots. Each of these collections reports changes for events that begin with a particular color (*e.g.,* green), and each of the three graphs in a cluster reports on changes from this initial color to another color (*e.g.,* green to yellow), or shows that the color did not change. The x-axis gives the scale factor that was applied to the ballistic coefficient as a proxy for the equivalent change in atmospheric density. Thus, these factors represent the atmospheric density error. Furthermore, each of these plots displays two lines: one showing the results for the cases in which the covariance is not altered to account for the expected atmospheric density error and another for the results when the (known) density errors are represented in the covariance through the use of a consider parameter. There are two morphological features common to the graphs regardless of starting color. Graphs that document the cases for which the same color remains after density alteration all show a value of 100% when the scale factor is unity. If the density is unaltered, the re-execution as expected reproduces the original result. A corresponding opposite to this is shown in graphs that document results for a color change. These must all show a 0% value for a scale factor of unity. Other than this, results behavior varies widely graph to graph and benefits from individual discussion. It should be kept in mind that the precise values of the numerical results below can, strictly speaking, be applied only to the data periods analyzed, but the general trends and relative performance are believed to be broadly representative.

6.1 "Green" Graph Set (Figure 2)

This results set documents false alarm (Type I) errors: events that are truly green but, due to density uncertainty, are mistakenly recategorized as either yellow or red. One notices that, in the main, there is rather little recategorization taking place here. Without error compensation in the covariance (consider parameter), less than 1% of the events are miscategorized. If compensation is included, the miscategorization rate rises to around 7%. From this results set alone, it would appear to be disadvantageous to include the characterized density error in the covariance. However, a final verdict on the utility of the practice must wait until performance against the other color-change categories is considered. The percentage of

recategorized events here is rather small, and the creation of Type I errors is far less concerning than the creation of Type II errors—some amount of labor may be wasted on such events, but at least truly worrisome events are not neglected due to miscategorization as benign ones.

6.2 "Yellow" Graph Set (Figure 3)

In this results set, larger levels of recategorization are observed, especially in the uncompensated case. The top graph shows that including the known density error in the covariance adds a substantial degree of stability to the situation: only 20 to 30% of the cases in which compensation is added (*i.e.,* the consider parameter applied) end up recategorized, whereas the uncompensated cases can recategorize up to 90% of the events, in both directions (yellow to green and yellow to red). The yellow to green (bottom) plot gives results for the Type II error of improperly demoting a yellow event to a green one, and it is clear that the consider parameter situation performs much more favorably. For the yellow to red (middle) case, the consider parameter situation performs slightly worse but at much less worrisome percentage levels than the uncompensated case. Use of the consider parameter emerges as the clear favorite here, producing a few more Type I errors but avoiding a large number of Type II errors. In both cases, however, the amount of recategorization argues for the need for general improvements in the fidelity of atmospheric density modeling.

6.3 "Red" Graph Set (Figure 4)

Here are reflected two strains of Type II errors: true red events that are miscategorized as yellow and true red events miscategorized as green. Both the compensated (consider parameter) and uncompensated cases reveal a disturbing amount of potential miscategorization as density errors are introduced. Examining the middle graph, one observes that for smaller density errors (scale factors from 0.7 to 1.25) both compensated and uncompensated approaches perform similarly (about 25% of red events improperly categorized as yellow). As the errors increase, the compensated case appears to perform more poorly. However, this apparent increase in poor performance over the uncompensated case is actually a mask: the blue line drops here because the uncompensated case is not miscategorizing so many red events as yellow but as green events—a far more serious Type II error. This dynamic is confirmed by examining the bottom plot, in which it is clear that miscategorization of red to green can reach as high as 80% for the uncompensated case while never venturing beyond 10% for the compensated (consider parameter) case. If one remains with the compensated situation, miscategorization rarely moves beyond a demotion from red to yellow (which means that the operational crews are actively monitoring the conjunction and are likely to see it move to red status as the propagation interval decreases). The uncompensated approach truly can set operational crews up for late-notice and perhaps even no-notice events, and no small number of them.

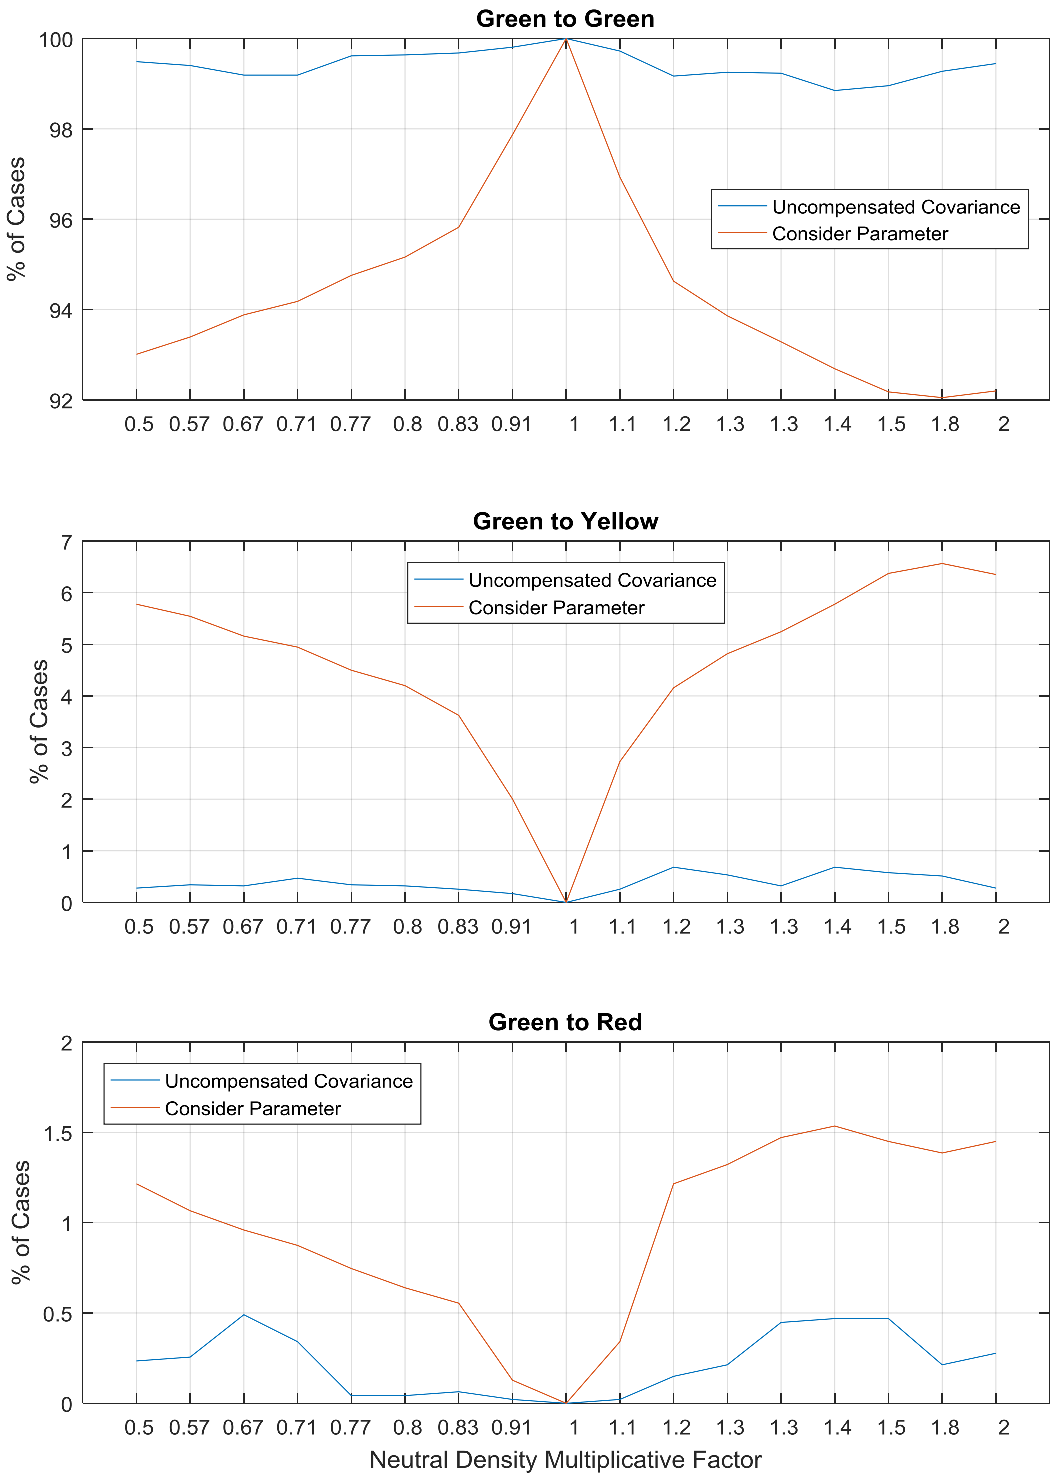

Figure. 2: Color Changes for Green Events as a Result of Density Error Addition

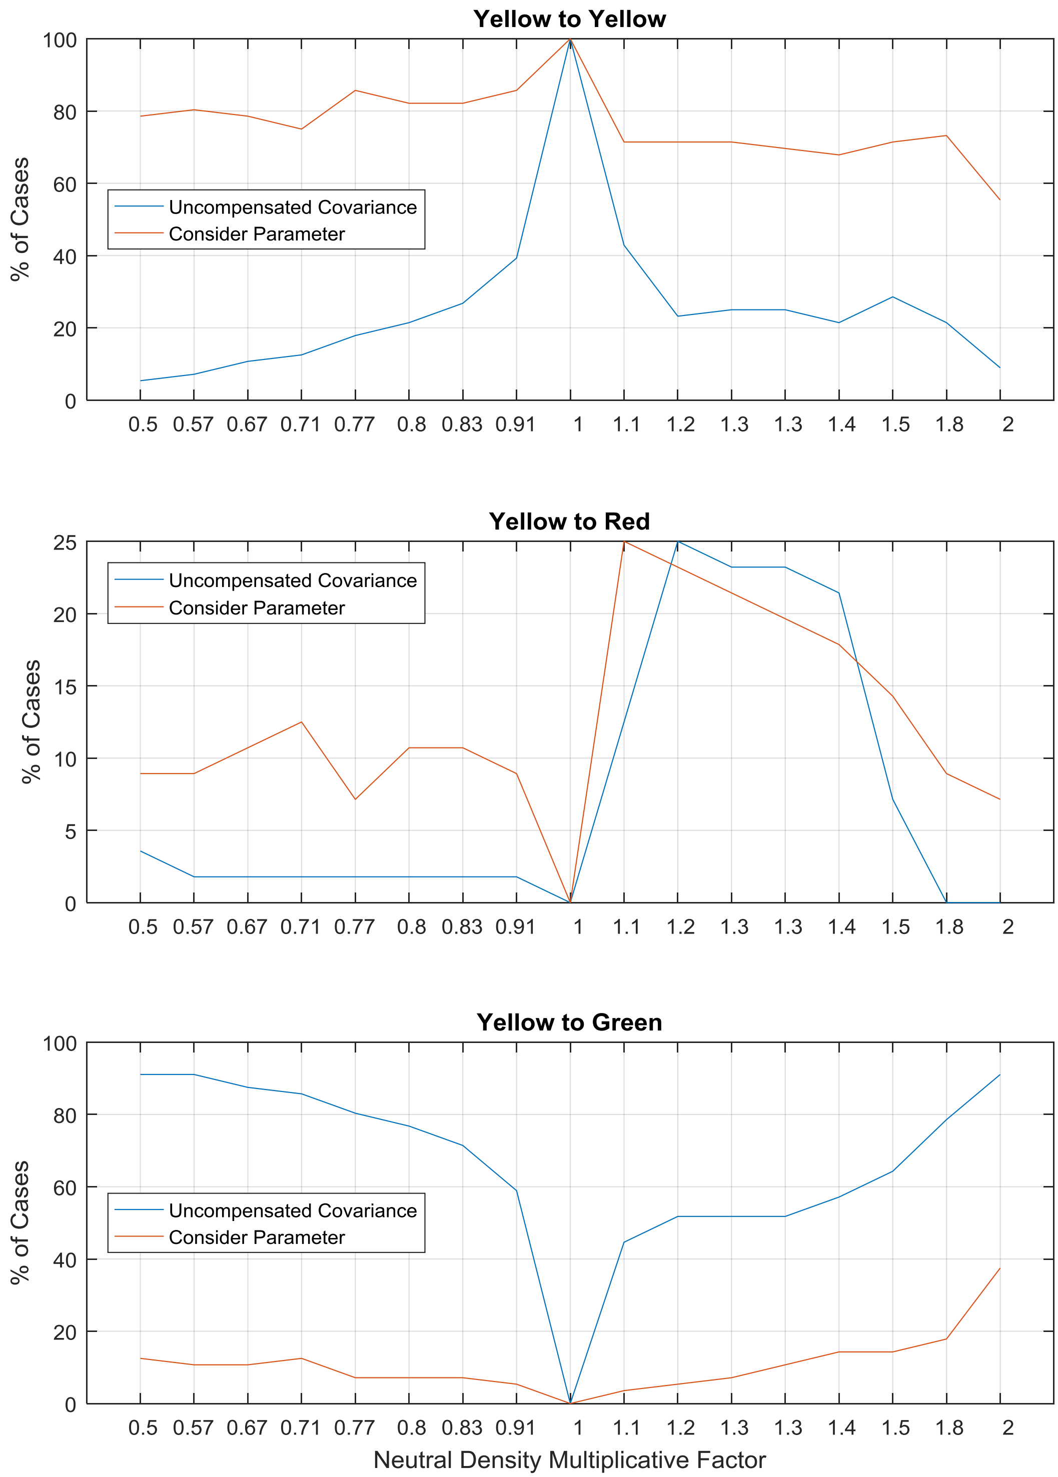

Figure. 3: Color Changes for Yellow Events as a Result of Density Error Addition

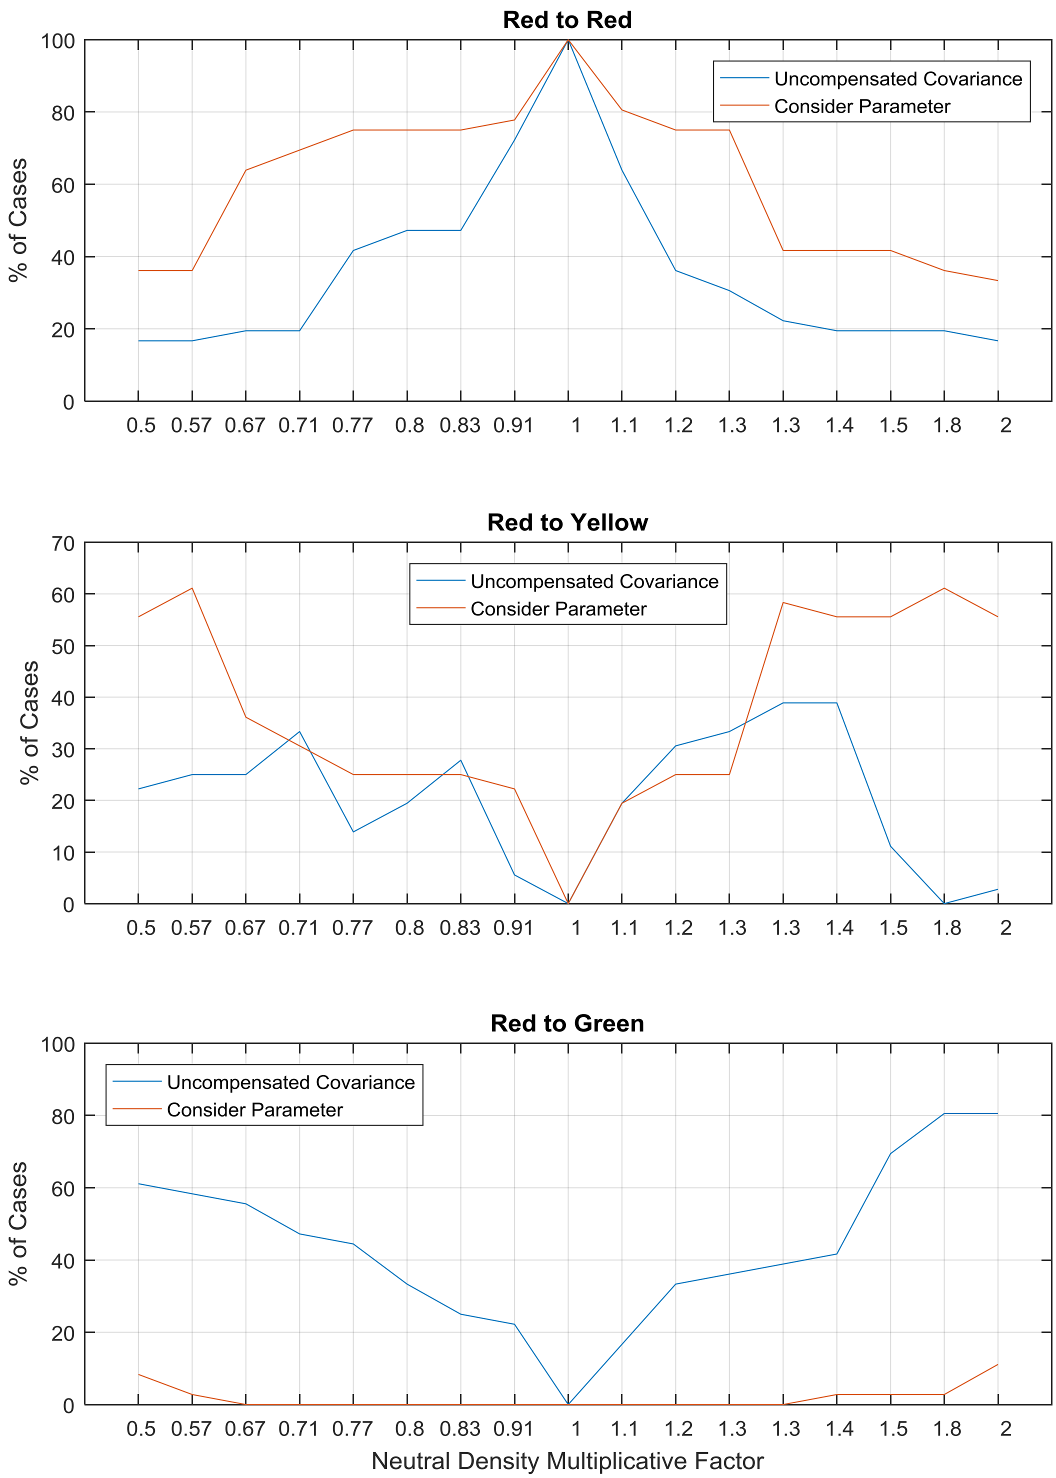

Figure. 4: Color Changes for Red Events as a Result of Density Error Addition

6.4 Interpretation of Results

It is expected that the specifics of the above results would be, to some degree, proper to the particular data period chosen to examine. Therefore, it is not the intent of the present analysis to advance actual precise levels of Type I and II errors expected at certain levels of density mismodeling. Rather, the general levels of event miscategorization, and the differences in these levels between the cases in which the covariance is and is not adjusted for the expected density error, are expected to be broad conclusions that can inform research objectives for atmospheric density prediction advances.

Surprisingly high levels of Type II error production occur in CA event categorization when even modest atmospheric density estimation errors enter into the calculation and are not considered as part of the conjunction satellites' state covariance matrices (in this analysis by adding a consider parameter). The Type II errors lead serious conjunction events to be dismissed as events requiring no subsequent analysis at all. Particularly worrisome are the "Yellow to Green" and "Red to Green" situations shown as the bottom plots of Figures 4 and 5: at certain density levels, over 80% of both potential and actual high-risk events will, due to miscategorization, be set aside completely by the risk assessment process, and for much of the dynamic range of the density errors, the Type II error rate exceeds 40%. Type I errors, while not of the same existential import as Type II, cause considerable additional operator workload for events that, as TCA approaches, will probably disappear. In the main, density error without covariance compensation generates rather small levels of Type I errors, although there are situations (e.g., overstatement of actual density for the yellow-to-red situation, middle graph of Figure 4) in which this error type becomes significant.

When characterization of the atmospheric density error is included in the objects' covariances, then the situation reverses. For events that are originally green, including the error compensation in the covariance can at worst produce about an 8% Type I error rate, whereas using the uncompensated covariance keeps this error rate at or below ~1%. For events that are originally yellow but after alteration are shown as red (the one type of Type I error for events that are yellow originally), for much of the space of induced density, error rides five to ten percentage points higher than that for the uncompensated covariances. While these Type I error rates are regrettable, given the substantial improvement in Type II error rates realized by active compensation of the covariance for expected density estimation error (as described in the above paragraph), they certainly constitute an acceptable operational trade-off, at least under present operational conditions. If the space catalogue were to increase in size substantially, as may occur when the US Air Force deploys their S-Band Radar, an 8% Type I error rate against Green events may no longer be sustainable with present operational practices. So even though the characterization of neutral atmospheric density estimation error and its incorporation in the covariance improves the operational situation substantially, it is not an outright substitute for an improvement in atmospheric density modeling. Both activities taken together are the only reliable approach to reducing both Type I and Type II CA risk assessment errors.

7 Conclusions and Future Work

The principal takeaway from the present experiment is that the ability to characterize the inherent error in neutral atmospheric density models, so that this error can be incorporated into space objects' state estimate covariances, can have a profound effect on mitigating the effects of such errors on CA risk assessment. While this practice does increase the Type I error rate

modestly, it can substantially improve the Type II error rate and prevent serious events from being dismissed outright. It instead preserves such events as yellow events—in a monitoring phase—so that as such events manifest their actual risk as TCA is approached, CA practitioners have already been made aware of them and have been preparing owner/operators for the possibility of such events becoming serious. Most atmospheric models presently in operational use do not include any embedded error estimation so that error information, tailored to the particulars of the situation, can be provided to users. This estimation should be a standard feature for future models and would be an extremely useful independent study and enhancement effort for presently-employed operational models.

The experimental results also show the overall value in reducing neutral atmospheric density estimation errors and, in particular, bringing them down to a bounded set of values about the true value. The larger error values explored in the experiment, which can represent commonly-encountered situations during space weather events such as coronal mass ejections, can cause quite serious miscategorizations of events even when compensation for these errors is included in the objects' state estimate covariances. Furthermore, the "Red-to-Yellow" Type II error (Figure 4, middle plot), while still addressable within current operational procedures, is not tenable with large space catalogues as there will be simply too many Yellow conjunctions to examine individually. A robust future for CA risk assessment, in which both larger space catalogues and the full range of space weather events can be addressed, requires methods that bring estimation errors down to relatively small and bounded values over the entire range of space weather conditions.

A further CA research activity that could potentially be helpful to atmospheric scientists would be to expand the present analysis to examine in greater detail the phenomenology of CA event miscategorization at smaller density errors. If a more detailed analysis were conducted in order to examine very closely the effects on CA of error values from, say, 0 to 30%, atmospheric researchers could have a better sense of the trade space between model accuracy and CA effects, which could help define performance targets for atmospheric density modeling improvement efforts.

8 References

Alfano, S. "Determining a Probability-Based Distance Threshold for Conjunction Screening." AAS Space Flight Mechanics Conference (Paper #13-352), Kauai HI, February 2013.

Alfano, S. "Relating Position Uncertainty to Maximum Conjunction Probability." *Journal of the Astronautical Sciences*, Vol. 53 No. 2 (April-June 2005), pp. 193-205.

Cerven, W.T. "Covariance Error Assessment, Correction, and Impact on Probability of Collision." 2011 AAS Space Flight Mechanics Conference (paper # 11-154), New Orleans LA, February 2011.

Cerven, W.T. "Improved Empirical Covariance Estimation." 2013 AAS/AIAA Astrodynamics Specialist Conference (paper # 13-768), Hilton Head SC, August 2013.

Chan, F.C. *Spacecraft Collision Probability*. El Segundo, CA: The Aerospace Press, 2008.

Duncan, M. and Long, A. "Realistic Covariance Prediction for the Earth Science Constellation." 2006 AAS/AIAA Astrodynamics Specialist Conference (paper # 06-6293), Keystone CO, August 2006.

Emmert, J., Byers, J., Warren, H., and Segerman, A. "Propagation of Forecast Errors from the Sun to LEO Trajectories: How Does Drag Uncertainty Affect Conjunction Frequency?" 2015 AMOS Technical Conference, Wailea HI, September 2015.

Emmert, J.T., Warren, H.P., Segerman, A.M., Byers, J.M., and Picone, J.M. "Propagation of Atmospheric Density Errors to Satellite Orbits." *Advances in Space Research* 59 (2017), pp. 147-165.

Foster, J.L. and Estes, H.S. "A Parametric Analysis of Orbital Debris Collision Probability and Maneuver Rate for Space Vehicles." NASA/JSC-25898 (August 1992).

Frisbee, J.H. "Re-examining Probability Dilution." 2009 AAS Astrodynamics Specialist Conference (Paper #09-413), Pittsburg PA, August 2009.

George, E.R. and Chan, F.K. "Covariance-Based Pre-Filters and Screening Criteria for Conjunction Analysis." 2012 AMOS Technical Conference, Wailea HI, September 2012.

Hejduk, M.D. and Johnson, L.C. "Approaches to Evaluating Probability of Collision Uncertainty." 2016 AAS Space Flight Mechanics Meeting (paper # 16-241), Napa CA, February 2016.

Hejduk, M.D., Plakalovic, D., Newman, L.K., Ollivierre, J.C., Hametz, M.E., Beaver, B.A., and Thompson, R.C. "Trajectory Error and Covariance Realism for Launch COLA Operations." 2013 AAS/AIAA Space Flight Mechanics Meeting (paper # 13-355), Kauai, HI. February 2013.

Hoots, F.R., Crawford, L.L, and Roehrich, R.L. "An Analytic Method to Determine Future Close Approaches between Satellites." *Celestial Mechanics* Vol. 33 No. 2 (June 1984), pp. 143-158.

Montenbruck, O. and Gill, E. *Satellite Orbits: Models, Methods, and Applications*. Berlin: Springer Verlag, 2005.

Narvet, S., Frigm, R.C., and Hejduk, M.D. "Assessment of Uncertainty-Based Screening Volumes for NASA Robotic LEO and GEO Conjunction Risk Assessment. 2011 AAS Astrodynamics Specialist Conference (Paper #11-432), Girdwood AK, August 2011.

Newman, L.K., Frigm, R.C., Duncan, M.G., and Hejduk, M.D. "Evolution and Implementation of the NASA Robotic conjunction Assessment Risk Analysis Concept of Operations." 2014 AMOS Technical Conference, Kihei HI, September 2014.

Poore, A.B., Ed. "Covariance and Uncertainty Realism in Space Surveillance and Tracking." Report of the Air Force Space Command Astrodynamics Innovation Committee, 27 JUN 2016.

Tapley, B.D., Schutz, B.E., and Born, G.H. *Statistical Orbit Determination*. New York: Elsevier Inc., 2004.

Vallado, D.A. and Finkleman, D. "A Critical Assessment of Satellite Drag and Atmospheric Density Modeling." *Acta Astronautica* 95 (2014), pp. 141-165.

Vallado, D.A., Hujsak, R.S., Johnson, T.M., *et al.* "Orbit Determination using ODTK Version 6." Fourth International Conference for Astrodynamic Tools and Techniques, European Space Astronomy Centre (ESA/ESAC), Madrid Spain, 2010.

Zaidi, W.H. and Hejduk, M.D. "Earth Observing System Covariance Realism." 2016 AAS/AIAA Astrodynamics Specialists Conference, Long Beach CA, September 2016.

9 Acknowledgements

Both authors have performed this research as part of their general duties in support of either NASA's Robotic Conjunction Assessment Risk Analysis program (Hejduk) or Headquarters Air Force Space Command's Studies and Analysis Division (Snow). Neither author presently holds any other professional affiliations, contractual or otherwise; and neither supports, or in the future expects to support, efforts to perform atmospheric modeling research and development.

The dataset used to perform the present analysis is a set of satellite conjunction screening results data from the NASA robotic conjunction assessment program archive. Such screening results provide satellite state and covariance information for two conjuncting satellites at TCA. These data originate from the USSTRATCOM Joint Space Operations Center and are a Controlled Unclassified Military Information (CUMI) datatype; as such, they are not publicly releasable and are exempt from Freedom of Information Act (FOIA) requests. Because of this article's general interest to the space weather community, it has been granted a special exception from the AGU's study data availability policy.